

Deciding when to stop: Efficient stopping of active learning guided drug-target prediction

Maja Temerinac-Ott^{1*}, Armaghan W. Naik², and Robert F. Murphy^{1,2,3}

¹ Freiburg Institute for Advanced Studies, University of Freiburg, Germany

² Computational Biology Department, Carnegie Mellon University, Pittsburgh, PA, USA

³ Departments of Biological Sciences, Biomedical Engineering and Machine Learning, Carnegie Mellon University, Pittsburgh, PA, USA
temerina@frias.uni-freiburg.de

Abstract. Active learning has shown to reduce the number of experiments needed to obtain high-confidence drug-target predictions. However, in order to actually save experiments using active learning, it is crucial to have a method to evaluate the quality of the current prediction and decide when to stop the experimentation process. Only by applying reliable stopping criteria to active learning, time and costs in the experimental process can be actually saved. We compute active learning traces on simulated drug-target matrices in order to learn a regression model for the accuracy of the active learner. By analyzing the performance of the regression model on simulated data, we design stopping criteria for previously unseen experimental matrices. We demonstrate on four previously characterized drug effect data sets that applying the stopping criteria can result in upto 40% savings of the total experiments for highly accurate predictions.

Keywords: active learning, drug-target prediction, simulation, matrix factorization, regression

1 Introduction

A critical step in developing new therapeutics is frequently to conduct large scale searches for potential drugs that can affect a desired target. Recently, it has become clear that finding successful drugs also requires searching for the absence of undesired effects on other targets. This need cannot be met by exhaustive experimentation, but selective experimentation driven by machine learning (a process referred to as *active learning*) may provide an alternative ([1]). The heart of active learning is having good predictive models to guide experimentation. Recent studies show that drug-target prediction algorithms can speed-up the discovery of new drugs (e.g., [2–4, 16]).

* This study was supported by BMBF e:Bio grant *Microsystems*, FKZ0316185. This paper was selected for oral presentation at RECOMB 2015 and an abstract is published in the conference proceedings.

Current drug-target prediction methods are coarse grained over at most a handful of 'campaigns'. In these, a classifier is trained with relatively large amounts of training data resulting from exhaustive screening, and then verified on a small test set. These data are generally identified manually, and limited to human 'expert' knowledge. This process is generally only performed once, or at most a handful of times due to the expense of exhaustive screening over many compounds. This procedure limits the generalization capability of the model and does not allow for an optimal exploration of the drug-target interaction space. Alternatively, active learning methods can be used to iteratively build a model of drug-target interactions. Instead of relying on large training data sets, the active learning procedure enlarges the training set stepwise, guided by the predictions on small, automatically-selected test sets. Thus time and costs are spent on improving the general model rather than having the verification of a small specific model that does not account for the large space of chemical compounds. The general model has the potential to predict side-effects early on in the drug design process, since a larger number of drugs are considered in the drug-target prediction matrix. A critical point when using active learning to guide experimentation is to decide when to stop, since the goal is to perform as few as possible experiments in order to have the best model. The best stopping time is reached when adding new experiments to the training set will not improve the accuracy on the test set. The difficulty, of course, is that calculating the true accuracy of the model requires all of the data. Therefore, reliable methods for predicting the accuracy of the current model during an active learning cycle are desired. Due to experimental cost and time restrictions, the best stopping time might not be desirable, so it would be helpful to stop earlier when a predefined confidence on the output of the model is reached.

Previous work in this area has generally addressed active learning methods or drug-target prediction methods, but rarely both. For example, active learning has been used to identify active compounds from a large pool of compounds targeting a single molecule [5]. Active learning has also been applied in the context of cancer research [24]. Several methods for drug-target prediction without active learning have been proposed recently [15, 6, 27, 20, 19, 18, 17] and remain an active area of research. The focus of this work is not to promote a particular drug-target prediction method, but to show using matrix factorization as an example how drug-target prediction can be combined with active learning and lead to reductions of experimentation cost. Initial results on applying active learning for drug-target prediction on multiple drugs and multiple targets simultaneously have been reported [7, 10], with and without requiring prior knowledge of drug or target similarities. In [10] the benefits of active learning on a large dataset from PubChem are reported, however without applying the stopping rule. In [7] an initial method for predicting the accuracy of active learning traces is presented, however it was not applied to the particular problem of drug-target prediction.

Several stopping rules for active learning have been considered in the past [22, 23, 26], however there has been little analysis of which performs the best in general. Four simple stopping criteria based on confidence estimation over

the unlabeled data pool and the label consistency between neighboring training rounds of active learning have been presented [26]. Instead of using a single criterion to stop, combining different stopping criteria in a feature vector describing the active learning trajectory has been proposed in [7]. The features of trajectories on simulated data are used to train a regression function in order to predict the accuracy of active learning algorithms on unseen simulated data. We will follow this approach and adopt it to the binary drug-target prediction case.

The major goals of our active learning system are: (1) We want to have a fast and reliable method to elucidate drug-target interactions. (2) Previous knowledge on similarities between drugs and similarities between targets should be included in the model, so that predictions for new drugs or targets (for which no experiments are available) are possible. (3) The number of experiments required to make confident predictions should be systematically reduced. (4) An efficient stopping rule for ending the active learning process should be designed.

Previously, kernel-based matrix factorization ([11]) has been shown to provide good models of drug-target interactions ([8]). In the kernelized Bayesian matrix factorization (KBMF) algorithm ([8, 9]), the drug-target interaction matrix is factorized by projecting the drugs and the targets into a common subspace, where the projected drug matrix and the projected target matrix can be multiplied in order to produce a prediction for the drug-target interaction matrix. The entries of the prediction matrix are modeled using truncated normal distributions. The projected drug matrix and target matrix are factored using two different kernels: a drug specific kernel and a target specific kernel. A kernel encodes the similarity between the drug and the target features. Thus prior information can be easily inserted to the model. Furthermore, the knowledge of the full interaction matrix is not needed in order to make predictions for new drugs, which is not the case for previous methods (i.e. [6]).

The main contributions of this work are: (i) We use KBMF to construct a powerful and practical active learning strategy for analyzing drug-target interactions. (ii) We extend previous work [7] on estimating the accuracy of active learning predictions to the KBMF case and show how it can be used to construct a stopping rule for experimentation. (iii) We provide a proof of concept through evaluation of the method on four data sets previously used for modeling of drug-target interactions [12]. (iv) We show the superiority of the proposed active learning approach compared to random choice of an equivalent number of experiments.

2 Methods

Active Learning Framework

An active learning method is an iterative process composed of four components: the initialization, the model, the active learning strategy and an accuracy measure for the predicted output in each step (Fig. 1). Most active learning papers focus on the second and third components. The active learning framework starts

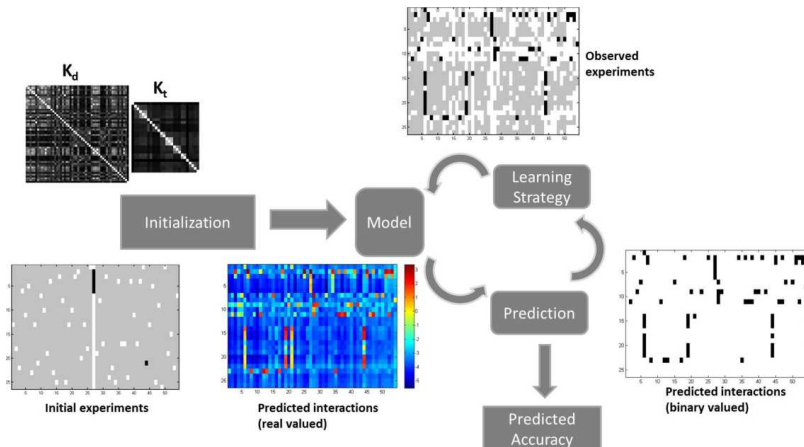


Fig. 1: The major components of an active learning framework. The entries of the matrix are color coded: label not known (light gray), interaction (black), no interaction (white). At initialization a subset of known labels for the interactions matrix and the drug and target kernels K_d and K_t are provided. In each round of the active learning algorithm, the labels of the entire interaction matrix are predicted and used to determine which labels to query next. In this figure, the dark red values represent a high probability for a hit, whereas the dark blue values represent a high probability for a miss.

with an initialization strategy which is followed by the generation of a model. The model is used to make predictions, in our application drug-target interactions are predicted. Interactions can be measured by performing an *experiment*, i.e. a direct assay of drug-target interaction (e.g., in cell extracts). Based on the predictions, an active learning strategy is applied to query new experiments (labels) which will improve the model. We use batchwise learning, where a fixed number of experiments is queried in each training round. Each training round defines a *time-point* in the active learning process and is measured by the number of batches of experiments performed. For each time-point the accuracy of the model is predicted. The process is stopped for example, if a certain budget for performing experiments is reached or the predicted accuracy of the model is high enough. We assume equal cost for each experiment.

Data Representation

We use interaction matrices $\mathbf{Y} \in \{-1, 1\}^{N \times M}$ to represent drug-target interactions. We assume that the outcome of the experiment determines the ground truth label $l \in \mathcal{L} = \{-1, 1\}$ for an interaction matrix entry. $N \in \mathbb{N}$ is the number of drugs, $M \in \mathbb{N}$ is the number of targets. Knowledge of the interaction between a drug $d \in \{1, 2, \dots, N\}$ and a target $t \in \{1, 2, \dots, M\}$ is ternary

encoded in the *experimental matrix* \mathbf{X} : +1 for an interaction, -1 for lack of interaction, and 0 to denote experiments which have not yet been performed. Hereby, the set of remaining experiments (unlabeled data) will be denoted by $\mathcal{X} = \{x = (d, t) | \mathbf{X}(x) = 0\}$. Therefore, we consider a semi-supervised binary labeling problem where the sign of the label indicates the interaction status between a drug and a target.

Kernelized Bayesian Matrix Factorization (KBMF)

As described previously [8, 9], KBMF can be effectively applied to model drug-target interactions. It approximates the interaction matrix by projecting the drug kernel $\mathbf{K}_d \in \mathbb{R}^{N \times N}$ and the target kernel $\mathbf{K}_t \in \mathbb{R}^{M \times M}$ into a common subspace of dimension $R \in \mathbb{N}$ such that the interaction matrix \mathbf{Y} can be reconstructed from the sign of its prediction matrix $\mathbf{F} \in \mathbb{R}^{M \times N}$:

$$\hat{\mathbf{Y}}(d, t) = \begin{cases} 1 & \text{if } \mathbf{F}(d, t) > 0 \\ -1 & \text{else.} \end{cases} \quad (1)$$

The prediction matrix \mathbf{F} is a product of the projected kernel matrices:

$$\mathbf{F} = ((\mathbf{A}_d)^T \mathbf{K}_d)^T ((\mathbf{A}_t)^T \mathbf{K}_t), \quad (2)$$

where $\mathbf{A}_d \in \mathbb{R}^{N \times R}$ and $\mathbf{A}_t \in \mathbb{R}^{M \times R}$ are subspace transformation matrices computed by the variational Bayes algorithm [8, 9] using the values of the experimental matrix \mathbf{X} . The dimension R of the subspace is a free parameter; we used the value of 20 previously determined to be optimal for these datasets [9]. The entries of the kernel matrix \mathbf{K}_d and \mathbf{K}_t are a measure of the pairwise similarities between drugs and targets respectively. The similarity matrices provided by Yamanishi et al. [12] and the KBMF implementation of semi-supervised classification¹ provided by Goenen [9] were used.

Initialization and experiment selection

Our initialization strategy is to select a random column and one random experiment from each row of the experimental matrix \mathbf{X} .

Uncertainty sampling We use uncertainty sampling ([13]) to form a batch of experiments $\{x_1, \dots, x_K\} \in \mathcal{X}$ by greedily choosing the $K \in \mathbb{N}$ experiments with the greatest uncertainty function U ([26]):

$$U(x) = - \sum_{l \in \mathcal{L}} P(l|x) \log P(l|x). \quad (3)$$

For the KBMF case the posterior probability is computed by the sigmoid function from the predicted interactions:

$$P(l = 1|x) = \frac{1}{1 + \exp(-\mathbf{F}(x))}, \quad (4)$$

¹ <http://research.ics.aalto.fi/mi/software/kbmf/>

and $P(l = -1|x) = 1 - P(l = 1|x)$ for no interaction respectively.

Stopping Rule

In order to stop the active learning process, a method is needed to predict the accuracy of the model for a given time-point along with the confidence of that prediction. As proposed previously in [7], the accuracy of a model at a given point in an active learning process can be predicted using a regression function trained for other, similar experimental spaces. The fully observed drug-target space is characterized by two measures, uniqueness (u) and responsiveness (r) [7] defined by:

$$r = \frac{1}{N \cdot M} \sum_{d,t,Y(d,t)=1} \mathbf{Y}(d,t) \quad (5)$$

$$u = \frac{uRows(\mathbf{Y}) + uColumns(\mathbf{Y})}{N + M}, \quad (6)$$

where $uRows(.)$ and $uColumns$ compute the number of unique rows and unique columns of a matrix.

The uniqueness and responsiveness are values in the range $[0, 1]$ and characterize the interaction matrix. Responsiveness measures the percentage of interactions in the matrix. Uniqueness is a measure of independence of the rows and columns in the matrix. The higher the value for uniqueness is, the more difficult it is to make predictions.

These two measures have two purposes: (1) They are used to compute features for a time-step in our current active learning process. (2) They can be used to generate simulation data having similar properties to the measured experimental data.

Each time-point t_i is described by a vector of 13 features $f_{t_i} \in \mathbb{R}^{13}$, defined as:

- $f(1), f(2)$: average observed responsiveness across columns (respectively rows)
- $f(3), f(4)$: average predicted responsiveness across columns (respectively rows)
- $f(5)$: average difference in predictions from last prediction for current time-point (t_i)
- $f(6)$: average difference in predictions from last prediction for previous time-point (t_{i-1})
- $f(7)$: fraction of predictions at t_{i-1} observed as responsive ($l = 1$) at t_i
- $f(8), \dots, f(10)$: minimum, maximum and mean number of experiments that have been performed for any drug
- $f(11), \dots, f(13)$: minimum, maximum and mean number of experiments that have been performed for any target

These features are normalized to the range $[0..1]$ and are further extended by computing the square root of their pairwise products.

To learn the accuracy predictor via simulation data, interaction matrices of size 50×50 were randomly sampled in the grid of uniqueness and responsiveness parameters 5%, 10%, ..., 95%. For each interaction matrix we derived 'perfect' Gaussian similarity kernels $\mathbf{K}_d, \mathbf{K}_t$ by pairwise distances of the column-space and row-space, respectively. These were disrupted by forcing 0%, 5%, 10% of the kernel entries to the value 1 and regularized to ensure positive semidefiniteness. Features computed from trajectories of the uncertainty sampling active learner on these data were collected; for each trajectory we also measured the accuracy of prediction against the ground truth. A linear model of these features against adjusted accuracies (accuracy above the fraction of experiments performed so far) was fitted by lasso regression ([14]). The lasso regularization parameter was chosen by 11-fold cross validation under squared loss, with holdout granularity at the level of trajectories. To make accuracy predictions from adjusted accuracy predictions, we added the fraction of experiments performed so far.

3 Results

For validation of our method, experiments are performed on four data sets extracted from the KEGG BRITE ([21]), BRENDA [25], SuperTarget [28] and DrugBank [29] databases, previously described by Yamanishi et al [12]². The data set consists of four drug-target interaction matrices: Nuclear Receptor, GPCR, Ion Channel and Enzyme.

Comparison of active and random learning strategies

In order to evaluate the efficiency of active learning methods, we compared the uncertainty sampling strategy with random choice of experiments in each time-step. On all four data sets the active learning strategy outperformed the random strategy (Fig. 2). On the GPCR and the Ion Channel dataset the active learning strategy using uncertainty sampling reaches 99% accuracy 5-6 times faster than the random strategy.

Predicting the accuracy of the model

As discussed in the introduction, in practice we require a mechanism to decide when to stop experimentation. It is not enough to have a good active learning method without the possibility to evaluate the accuracy of the whole model apart from acquiring all the data. Therefore we have simulated interaction matrices with uniqueness, responsiveness in the range [0.05..0.95] and kernel noise in the range [0..0.1]. We then performed active learning simulations using our KBMF model and uncertainty sampling and learned a regression function for the predicted accuracy. The results of applying the regression function to the computed features at each time point are shown in red in Fig. 3 for the four data sets. On

² <http://web.kuicr.kyoto-u.ac.jp/supp/yoshi/drugtarget/>




Fig. 2: Comparison of random sampling (green) to uncertainty sampling (black) on the four data sets. The solid and the dotted line represent respectively the mean and the standard deviation of 5 random initializations. For random sampling five random runs were performed for each initialization.

all four data sets, the predicted accuracy of 90% guarantees the true accuracy to be at least 90%, and the predicted accuracies are a reasonable lower estimate for the true accuracy.

Learning the stopping rule

Statistics on the performance of the accuracy predictor in simulations can be used to design a stopping rule [7]. We adopt this method to determine a threshold for stopping the active learning procedure. The simulated data is used to assess the probability that the true accuracy is greater than or equal to the predicted accuracy using 11-fold cross-validation. We count for each predicted accuracy value how often the condition was fulfilled and divide it by the total occurrence of this predicted value (Fig. 4). As expected, a low predicted accuracy with a value below 0.5 will have a high probability to be measured higher in the actual experiment. Predicted values below 0.5 are not of interest, since the predicted value is too low. In the beginning of the active learning procedure a small amount of data is available, so it is hard to make good predictions about




Fig. 3: The true accuracy (black) and the predicted accuracy (red) are shown for the four data sets.

the accuracy of the method. However, the more data is gathered in the active learning procedure, the more confident the predictor gets, reaching a peak for predicting the accuracy of 0.8 and higher for 65% of the cases. For very high accuracies (> 0.95), the chance that the actual accuracy exceeds the prediction naturally drops drastically. From Fig. 4 the best threshold to stop lies in the range 0.8 to 0.9. Since higher accuracy values are more desirable, our stopping rule was to terminate the active learning procedure, when the predicted accuracy is 0.9.

Applying the stopping rule

In the original KBMF paper ([8]), the KBMF classifier was evaluated by 5-fold five cross validation using 80% of drugs for training and 20% of drugs for testing. We wanted to test if a matching accuracy on the test set could be reached by choosing actively a reduced number of experiments for training. In other words, assuming that we get to perform selected experiments drawn from a given set, use them to train a model, and make predictions for a held out set (for which experiments are not possible), can we get an accurate model without doing all experiments? For this purpose our active learning strategy was modified. We use




Fig. 4: The probability that a predicted accuracy is below or equal to the true accuracy is plotted against the threshold.

Table 1: Average AUC on hold out data and percentage of experiments after applying our stopping rule. The average AUC obtained on held out data using 80% of the data for training [8] is compared with the average AUC obtained by training with only the listed percentage of experiments obtained by applying the stopping rule. The percentage of experiments can be halved by using the proposed stopping rule.

Dataset	Goenen results AUC (%)	With stopping rule	
		AUC(%)	experiments (%)
Nuclear Receptor	82.4	81.7	52.9
GPCR	85.7	81.6	39.3
Ion Channel	79.9	83.8	44.2
Enzyme	83.2	77.8	29.7

1% of drugs as the batch size and select in each run the drugs which the classifier is the most uncertain about. For uncertainty sampling using the predictions of the KBMF classifier, this means that drugs with the maximal mean uncertainty across targets are selected.

Instead of using all 80% of drugs for training ([8]), we use the predicted accuracy on the training data to stop acquisition. When the predicted accuracy on the training set reaches a threshold, the active learning process is stopped and the AUC value on the test set (the 20% of the drugs which were held out) is reported. The average results after 5-fold five cross validation are reported in Table 1. By using the stopping-rule on all four data sets, only half of the drugs were needed for training to reach a similar AUC value to that when using all 80% of the drugs for training.

Table 2: Average difference between the BST point and the stopping point chosen by various stopping rules, over all evaluation data sets. OU=Overall Average Uncertainty, MEE=Minimum Expected Error PA=Predicted Accuracy. The value in the brackets denotes the threshold. The smaller the difference Δ_{ave} value is, the better the stopping criterion is.

Methods	OU(0.12)	OU(0.09)	OU(0.06)	OU(0.03)	OU(adapted)
$\Delta_{ave}(\%)$	40.1 (± 12.2)	33.8 (± 17.8)	40.1 (± 21.3)	50.9 (± 5.4)	28.2 (± 29.1)
Methods	MEE(0.12)	MEE(0.09)	MEE(0.06)	MEE(0.03)	MEE(adapted)
$\Delta_{ave}(\%)$	40.1 (± 11.7)	38.3 (± 12.7)	36.1 (± 13.4)	40.6 (± 12.1)	30.3 (± 12.6)
Methods	PA(0.85)	PA(0.9)	PA(0.95)		
$\Delta_{ave}(\%)$	32.8 (± 8.8)	13.7 (± 11.3)	22.1 (± 15.4)		

Comparison Of Stopping Rules

We compared the stopping criteria overall uncertainty (OU) and minimum expected error (MEE) with a fixed threshold as well as an adapted threshold based on label consistency as described [26] with the stopping method based on predicted accuracy [7] in Table 2. As in [26] we use the absolute difference of the percentage of experiments completed at the stopping time-point to the percentage of experiments completed at the best stopping time (BST) averaged over four data sets (Δ_{ave}) to evaluate different stopping criteria. In [26] the BST is defined as the time-point, when the classifier first reaches the highest performance. The predicted accuracy (PA) method with threshold 0.9 produces the smallest average error to the BST. Both MEE and OU perform two to three times worse than the PA method, even with the adaptive threshold method. The fixed threshold for OU and MEE fails on average, because each of the four data sets has a different optimal threshold for OU and MEE. The maximum uncertainty (MU) and the selected accuracy (SA) stopping criteria [26] could not be applied, since those curves are not continuous on these data sets.

4 Conclusions and Discussion

We have presented an active learning method for prediction of drug-target interactions based on kernelized matrix factorization. Building on prior work [8], our model can efficiently leverage prior information through kernels to achieve high predictive accuracy. We have furthermore shown that our method can significantly improve the prediction task for drug-target interactions when only a limited number of experiments can be performed. For three real-world data sets with high uniqueness values, the active learning strategy achieves 99% accuracy with 2-3 times fewer experiments than a random sampling strategy. It is important to note that our goal was not to choose the best possible matrix completion method for these specific datasets, but to show that a good method can be used as a basis for active learning to dramatically reduce further experimentation.

Please note, that the presented framework is not limited to KBMF only. Any other model for drug target prediction could be applied that produces outputs for drug-target scores which can be converted into probabilities. Furthermore the selection strategy uncertainty sampling could be replaced by other active learning strategies (i.e. diversity sampling) to learn new traces on simulated data. The active learning features could be improved by feature selection methods.

For a practitioner to realize these advantages, we have provided a method for estimating the accuracy of an actively learned model using only experimental results already collected; this estimated accuracy is generally a lower bound of the true accuracy of the model. We have shown that this method, calibrated from simulation data, accurately assesses the active learner performance on our real-world data. We have also shown that by applying a stopping rule learned on the simulated data, only half of the experiments are needed to achieve similar accuracies on holdout data. We conclude that active learning driven experimentation is a practical solution to large experimental problems in which time or expense make exhaustive experimentation undesirable.

References

1. Murphy, R.F.: An active role for machine learning in drug development. *Nature Chemical Biology* 7, 327–330 (2011)
2. Besnard, J., Ruda, G.F., Setola, V., Abecassis, K., Rodriguiz, R.M., Huang, X., Norval, S., Sassano, M.F., Shin, A.I., Webster, L.A., Simeons, F.R.C., Stojanovski, L., Prat, A., Seidah, N.G., Constam, D.B., Bickerton, G.R., Read, K.D., Wetsel, W. C., Gilbert, I.H., Roth, B.L., Hopkins, A.L.: Automated design of ligands to polypharmacological profiles. *Nature* 492, 215–220 (2012)
3. Paolini, G.V., Shapland, R.H.B., van Hoorn, W.P., Mason, J.S., Hopkins, A.L.: Global mapping of pharmacological space. *Nature Biotechnology* 24, 805–815 (2006)
4. Reymond, J.-L., van Deursen, R., Blum, L.C., Ruddigkeit, L.: Chemical space as a source for new drugs. *MedChemComm* 1, 30–38 (2010)
5. Warmuth, M.K., Liao, J., Rätsch, G., Mathieson, M., Putta, S., Lemmen, C.: Active Learning with Support Vector Machines in the Drug Discovery Process. *J. Chem. Inf. Comput. Sci.* 43, 667–673 (2003)
6. Atias, N., Sharan, R.: An Algorithmic Framework for Predicting Side Effects of Drugs. *Journal of Computational Biology* 18(3): 207–218 (2011)
7. Naik, A.W., Kangas, J.D., Langmead, C.J., Murphy, R.F.: Efficient modeling and active learning discovery of biological responses. *PLoS ONE* 8(12), 83996 (2013)
8. Gönen, M.: Predicting drugtarget interactions from chemical and genomic kernels using Bayesian matrix factorization. *Bioinformatics* 28, 2304–2310 (2012)
9. Gönen, M., Khan, S.A., Kaski, S.: Kernelized bayesian matrix factorization. In: *International Conference on Machine Learning*, pp. 864–872 (2013)
10. Kangas, J.D., Naik, A.W., Murphy, R.F.: Efficient discovery of responses of proteins to compounds using active learning (15), 143 (2014). *BMC Bioinformatics*
11. Bazerque, J.A., Giannakis, G.B.: Nonparametric basis pursuit via sparse kernel-based Learning. *IEEE Signal Processing Magazine* 30, 112–125 (2013)
12. Yamanishi, Y., Araki, M., Gutteridge, A., Honda, W., Kanehisa, M.: Prediction of drugtarget interaction networks from the integration of chemical and genomic spaces. *Bioinformatics* (24), i232–i240 (2008)

13. Lewis, D.D., Gale, W.A.: A sequential algorithm for training text classifiers. In: Proceedings of the eleventh Annual ACM-SIGR Conference on Research and Development in Information Retrieval, 3–12 (1994)
14. Tibshirani, R.: Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B*, Vol 58, No. 1, pp. 267–288 (1996)
15. Yamanishi, Y., Araki, M., Gutteridge, A., Honda, W., Kanehisa, M.: Prediction of drug-target interaction networks from the integration of chemical and genomic spaces. *Bioinformatics* 24, 232240 (2008)
16. Keiser, M.J., Setola, V., Irwin, J.J., Laggner, C., Abbas, A.I., Hufeisen, S.J., Jensen, N.H., Kuijter, M.B., Matos, R.C., Tran, T.B., Whaley, R., Glennon, R.A., Hert, J., Thomas, K.L.H., Edwards, D.D., Shoichet, B.K., Roth, B.L.: Predicting new molecular targets for known drugs. *Nature* 462, 17581 (2009)
17. Bleakley, K., Yamanishi, Y.: Supervised prediction of drug-target interactions using bipartite local models. *Bioinformatics* 25(18), 23972403 (2009).
18. Zheng, X., Ding, H., Mamitsuka, H., Zhu, S.: Collaborative matrix factorization with multiple similarities for predicting drug-target interactions. In: KDD, pp. 10251033 (2013)
19. Cheng, F., Liu, C., Jiang, J., Lu, W., Li, W., Liu, G., Zhou, W., Huang, J., Tang, Y.: Prediction of drug-target interactions and drug repositioning via network-based inference. *PlosComputational Biology* (2012)
20. Alaimo, S., Pulvirenti, A., Giugno, R., Ferro, A.: Drug-target interaction prediction through domain-tuned network-based inference. *Bioinformatics* (2013)
21. Kanehisa, M., Goto, S., Hattori, M., Aoki-Kinoshita, K.F., Itoh, M., Kawashima, S., Katayama, T., Araki, M., Hirakawa, M.: From genomics to chemical genomics: new developments in kegg. *Nucleic Acids Res* 34, 354357 (2006)
22. Laws, F., Schätze, H.: Stopping criteria for active learning of named entity recognition. In: Proceedings of the 22Nd International Conference on Computational Linguistics - Volume 1. COLING 08, pp. 465472. Association for Computational Linguistics, Stroudsburg, PA, USA (2008).
23. Vlachos, A.: A stopping criterion for active learning. *Computer, Speech and Language* 22(3), 295312 (2008)
24. Danziger, S.A., Zeng, J., Wang, Y., Brachmann, R.K., Lathrop, R.H.: Choosing where to look next in a mutation sequence space: Active learning of informative p53 cancer rescue mutants. *Bioinformatics* 23(13), 104114 (2007)
25. Schomburg, I., Chang, A., Ebeling, C., Gremse, M., Heldt, C., Huhn, G., Schomburg, D.: Brenda, the enzyme database: Updates and major new developments. *Nucleic Acid Research* 32, 431433 (2004)
26. Zhu, J., Wang, H., Hovy, E., Ma, M.: Confidence-based stopping criteria for active learning for data annotation. *ACM Trans. Speech Lang. Process.* 6(3), 31324 (2010)
27. Campillos, M., Kuhn, M., Gavin, A.-C., Jensen, L.J., Bork, P.: Drug target identification using side-effect similarity. *Science* 321, 263266 (2008)
28. Günther, S., Kuhn, M., Dunkel, M., Campillos, M., Senger, C., Petsalaki, E., Ahmed, J., Urdiales, E.G., Gewiess, A., Jensen, L.J., Schneider, R., Skoblo, R., Russell, R.B., Bourne, P.E., Bork, P., Preissner, R.: Supertarget and matador: Resources for exploring drug-target relationships. *Nucleic Acid Research* 36, 91922 (2008)
29. Wishart, D.S., Knox, C., Guo, A.C., Cheng, D., Shrivastava, S., Tzur, D., Gautam, B., Hassanali, M.: Drugbank: a knowledgebase for drugs, drug actions and drug targets. *Nucleic Acid Research* 36, 9016 (2008)